arXiv:1512.02306v1 [stat.AP] 8 Dec 2015

Nonparametric Reduced-Rank Regression for Multi-SNP, Multi-Trait Association Mapping

Ashlee Valente[1], Geoffrey Ginsburg[1], Barbara E Engelhardt[2]

1 Center for Applied Genomics and Precision Medicine, Duke University, Durham, NC 27708, USA 2 Department of Computer Science and Center for Statistics and Machine Learning, Princeton University, Princeton, NJ, USA

∗ **Corresponding author e-mail: bee@princeton.edu**

Abstract

Genome-wide association studies have proven to be essential for understanding the genetic basis of disease. However, many complex traits—personality traits, facial features, disease subtyping—are inherently high-dimensional, impeding simple approaches to association mapping. We developed a nonparametric Bayesian reduced rank regression model for multi-SNP, multi-trait association mapping that does not require the rank of the linear subspace to be specified. We show in simulations and real data that our model shares strength over SNPs and over correlated traits, improving statistical power to identify genetic associations with an interpretable, SNP-supervised low-dimensional linear projection of the high-dimensional phenotype. On the HapMap phase 3 gene expression QTL study data, we identify pleiotropic expression QTLs that classical univariate tests are underpowered to find and that two step approaches cannot recover. Our Python software, BERRRI, is publicly available at GitHub: `https://github.com/ashlee1031/BERRRI`

Introduction

Genetic variation is an important driver of variability in complex traits, contributing to phenotypic differences among individuals and across populations. To find genetic variants that regulate complex heritable phenotypes, genome-wide association studies (GWAS) have been used to identify associations between single nucleotide polymorphisms (SNPs) and quantitative traits such as height, weight, or blood pressure.

The statistical methods used for association mapping with organismal quantitative traits can also be used to identify genetic variation affecting cellular traits such as RNA expression levels or collections of

co-regulated or correlated transcripts (expression QTLs or eQTLs) [7, 63, 10]. In this framework, the quantitative traits are the expression of RNA transcripts. eQTL analyses have been used to suggest possible biological mechanisms for trait associated SNPs (TASs) identified using GWAS, because of evidence that genetic regulation of gene expression is a primary mechanism of the genetic regulation of disease risk [38, 21]. The availability of eQTL data for the target tissue of a complex disease provides insight into possible biological basis for disease risk and can help identify networks of genes involved in disease pathogenesis [12, 19, 32].

Typically, eQTLs are classified as cis- or trans-acting, meaning the regulatory target of the QTL is allele-specific or not, respectively. Gene regulatory networks suggest that trans-eQTLs might be due to cis-eQTLs regulating RNA that, in turn, affect gene expression of a distal gene in a dosage-specific way [39]. A polymorphism affecting a regulatory gene in cis could thus have observable downstream effects on the expression of its targets, resulting in *pleiotropic* cis- and trans-effects. Current methods identify pleiotropic eQTLs in one of three ways: i) identifying co-expressed transcripts [61, 39]; ii) using cis-eQTLs to identify trans-eQTLs [22]; or iii) projecting gene expression data down to a low dimensional space and performing association mapping in this space [23, 24]. In this paper, we perform simultaneous supervised projection of both the SNPs and the high-dimensional phenotype and association mapping using reduced rank regression.

Statistical challenges of eQTL analyses

Conventional eQTL analyses identify SNPs that explain a significant fraction of the variation in gene expression levels by testing for association of expression levels of a single gene with a single SNP; in these studies, gene expression levels are treated as quantitative traits [12, 37]. The univariate approach to association mapping in eQTL studies is plagued by limited statistical power to detect associations with small effects due to the large number of tests relative to the small sample sizes in these studies [45].

Cis-eQTLs have been more successfully identified than trans-eQTLs [29]; practically, cis-eQTLs are assumed to lie within 200Kb of the transcription start or end site (TSS,TES) of a gene, reducing the number of tests for association by three orders of magnitude [40]. In contrast, interrogating the entire genome for regulatory associations for each gene—as is done for trans-eQTLs—incurs a substantial multiple hypothesis testing penalty. It is unclear if the overrepresentation of discovered cis-eQTLs is due to lower statistical power to identify trans-eQTLs, or if this overrepresentation reflects different relative

numbers and effect size distributions across cis- and trans-eQTLs [59, 58, 37, 11]. However, recent studies have shown that, with sufficient sample sizes, hundreds of trans-eQTLs may be found and replicated within a tissue type [18, 30].

Multi-trait association mapping

One approach to mitigate problems resulting from limited sample size is to perform association mapping across multiple traits simultaneously [46, 16, 17, 41]. The idea behind multi-trait association mapping is that the statistical tests across multiple correlated traits (such as coronary heart disease and myocardial infarction, or gene transcripts with correlated expression levels) with shared underlying genetic associations will share strength across traits to improve statistical power [43]. These approaches have been effective at identifying small signals present over various measured traits. For example, the Genotype-Tissue Expression (GTEx) study data includes gene expression levels and genotypes for a large number of tissue samples from hundreds of individuals [3]. It is often the case that a cis-eQTL exists in multiple tissues. In GTEx, sharing strength across many tissues identifies many more ubiquitous cis-eQTLs, or cis-eQTLs that are shared across tissues, than univariate methods alone [16]. This approach may also be applied to correlated phenotypes [43], medical imaging traits such as fMRI data [53], and many other high-dimensional biological and clinical traits.

Multi-SNP analyses identify a subset of local SNPs as cis-eQTLs using various types of model selection methods [15, 57, 41, 31, 62]. These approaches evaluate associations with all SNPs simultaneously, often using a form of multivariate regression that includes sparsity-inducing priors on the effect sizes [50, 33]. The many correlated SNPs (due to linkage disequilibrium) included in these models relative to sample size makes classical sparse regression techniques like Lasso non-robust [31]; model averaging has been used used to create a robust multi-SNP mapping method [52].

Reduced Rank Regression (RRR)

In regression models for GWAS or eQTL analyses, when the number of predictors is much greater than the number of observations (the $p \gg n$ scenario), accurately estimating the regression coefficients is a challenge [34, 17, 56]. An alternative approach to inducing sparsity in the effect sizes is to project these high-dimensional coefficients down to a lower-dimensional basis, making estimation easier because of the many fewer statistical tests [2, 34, 17].

Reduced rank regression (RRR) [1, 2] does this by explicitly mapping the coefficients of a linear regression model down to a small number of latent factors. These multivariate regression models have been extended to multi-trait RRR models, which capture linear relationships between multiple response variables and multiple predictor variables on a low-dimensional linear manifold, sharing signal strength across co-regulated response variables. Inducing sparsity on that manifold further reduces the number of model parameters from standard multi-SNP, multi-trait regression models and leads to interpretable results [9, 17].

Canonical correlation analysis (CCA) also achieves this projection onto a lower-dimensional space to examine relationships between two high-dimensional data sets [62]. However, such approaches are best suited for exploratory data analysis as they do not lead to a measure of effect size for associations in the data or a posterior probability of SNP association [15].

RRR models have previously been developed for and applied to genomic study data. For example, a sparse RRR method was developed to identify predictive signatures of disease status in fMRI images from Alzheimer's disease populations [54, 53]. A second RRR method was developed for genetic association studies [25, 35], and applied to data from the Northern Finland Birth Cohort to uncover two novel loci associated with lipoprotein profiles [35].

Approach

Here we describe our model, Bayesian extendable RRR with an Indian Buffet Process prior (BERRRI), which goes beyond previous RRR approaches in three ways to enable genome-scale association mapping. First, our model uses a nonparametric prior on the reduced rank parameters, allowing the rank to be estimated from the data. Second, BERRRI uses binary values to indicate SNP inclusion in the low-dimensional subspace, leading to interpretable posterior probabilities of association for every SNP. Third, BERRRI uses a variational Bayes approach, making posterior estimates computationally tractable for large genomic studies.

Bayesian Extendable RRR Model

We propose a reduced rank regression model to quantify associations between SNPs and high-dimensional, correlated traits where the response matrix $Y \in \Re^{N \times P}$ is a set of P traits for N individuals. The genotype

data, $X \in \Re^{N \times Q}$ for N individuals and Q SNPs, are encoded by $x_{n,q} \in \{0, 1, 2\}$ representing the number of copies of the minor allele in individual n and SNP q. We assume that a weighted linear combination of the genotype data X approximates the trait matrix Y with Gaussian noise $\epsilon \in \Re^{N \times P}$:

$$Y = XZA + \epsilon \tag{1}$$

$$\epsilon_{n,\cdot} \sim N(0, \sigma_n^2 I) \tag{2}$$

As in a general reduced rank regression framework, we decompose the regression coefficients into a product of two matrices. Binary matrix $Z \in \{0, 1\}^{Q \times K}$ selects sparse sets of SNPs for each of K factors, and continuous matrix $A \in \Re^{K \times P}$ identifies subsets of traits that are associated with the SNPs included in one of K factors. Our model diverges from the canonical RRR model in the prior distributions on the A and Z matrices. In particular, matrix Z is a draw from the Indian Buffet Process (IBP) [28, 49]:

$$z_{q,k} \sim Bernoulli(\pi_k) \tag{3}$$

$$\pi_k \sim Beta\left(\frac{\alpha}{K}, 1\right) \tag{4}$$

$$A_{k\cdot} \sim \mathcal{N}(0, \delta_k . I). \tag{5}$$

K is inferred from the data, and is the number of nonzero columns in Z.

Moreover, we impose sparsity on the effect sizes, A, by placing a conjugate inverse-gamma prior on each parameter δ, known as an automatic relevance determination (ARD) prior [36, 51]. This has the effect of placing an ℓ_1-type penalty on the effect sizes (Figure 1).

The IBP is a Bayesian nonparametric prior on binary matrices with support on an infinite number of columns [28, 49]. In the context of RRR, the IBP allows us to infer i) which SNPs are predictive of the observed data and ii) how many latent components K are recovered (where each latent component is a separate multi-SNP, multi-trait association). The IBP prior has the desirable behavior that, in the posterior, we recover a finite number of components conditioned on a finite data set, and additional components may be added as more samples are collected to improve statistical power. Because Z is binary, we may interpret the elements in a model selection context: the posterior probability of $z_{q,k}$ represents the probability that SNP q is associated with the traits included in the kth component.

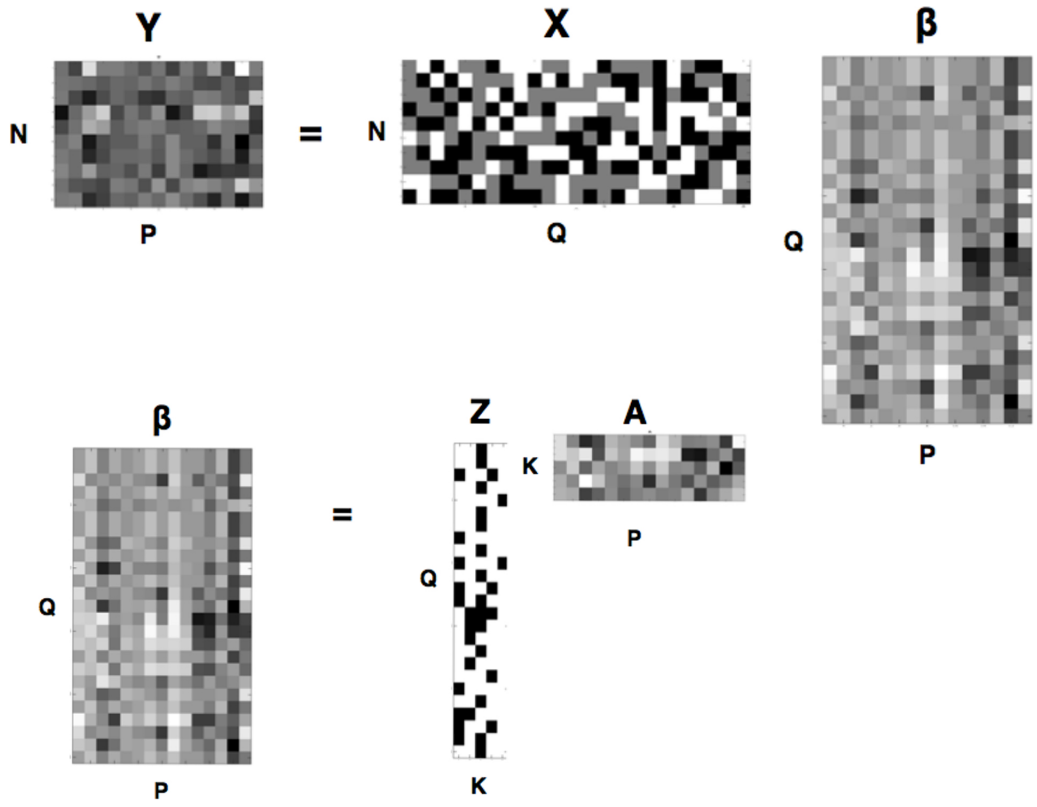

Figure 1: A visual representation of the BERRRI model. Y is a matrix of P correlated quantitative traits, X is a genotype matrix $x_{n,q} \in 0, 1, 2$, β is the matrix of effects, decomposed into Z, the binary matrix, and A, the matrix of continuous effect sizes.

IBP estimation is a computationally difficulty problem, with an enormous number—$2^{q \times K}$—of binary configurations. Sampling methods for posterior estimation in the IBP are computationally intractable and sensitive to initialization [13]. These problems are magnified in high-dimensional genomic data. A variational Bayesian approximate method was recently developed for the IBP [6, 13]. This variational approach exploits the stick breaking representation of the IBP [47] for fast and robust posterior estimation. We base posterior inference in the BERRRI model on this variational approach.

Variational Bayes and the IBP

Variational Bayes is an alternative to Monte Carlo sampling methods that provides an analytic solution to an approximation of the posterior distribution [4]. In contrast, sampling techniques provide an approximation to the exact posterior distribution. In many cases, variational Bayes methods provide solutions with greater computational efficiency than sampling methods, with minimal sacrifice in accuracy [48].

In the BERRRI model, the set of parameters to estimate are $W = A, Z, \pi, \delta$, and the fixed hyperparameters are $\theta = \alpha, \sigma_n^2$. We write the log probability of W:

$$\log p(W|Y, X, \theta) = \log p(W, Y|X, \theta) - \log p(Y|X, \theta), \tag{6}$$

where the log marginal probability:

$$\log p(Y|X, \theta) = \log \int p(Y, W|X, \theta) dW \tag{7}$$

is intractable to compute.

Mean field variational Bayes methods approximate this log marginal probability with a variational distribution $q(W)$ [5, 55]. We follow prior work [13] and use a finite IBP variational approach to these updates. Variational inference allows us to approximate model parameters λ, η, ϕ^m by minimizing the Kullback-Leibler divergence between an approximate variational distribution and a finite approximation to the IBP [13]. In our case, we use the following factorized variational distribution:

$$q(W) = q_{\lambda_k}(\pi_k) q_{\eta_{qk}}(Z_{qk}) q_{\phi_{kp}, \varphi_{kp}}(A_{kp}) q_\kappa(\delta) \tag{8}$$

where

$$q_{\lambda_k}(\pi_k) = Beta(\pi_k; \lambda_{k,1}, \lambda_{k,2}), \tag{9}$$

$$q_{\eta_{q,k}}(Z_{q,k}) = Bernoulli(Z_{q,k}; \eta_{qk}) \tag{10}$$

$$q_{\phi_{k,p},\varphi_{k,p}}(A_{k,p}) = \mathcal{N}(A_{k,p}; \phi_{k,p}, \varphi_{k,p}) \tag{11}$$

$$q_{\kappa_{k,p}}(\delta_{k,p}) = Gamma^{-1}(\delta_{k,p}; \kappa_{k,p,1}, \kappa_{k,p,2}) \tag{12}$$

We then optimize the variational parameters to minimize the Kullback-Leibler divergence $D(q \mid\mid p)$ to the finite approximate IBP model distribution p. This reduces to an optimization problem of maximizing the lower bound on $p(Y|X, \theta)$. While this optimization can be difficult, when the distributions are in the exponential family there exists a closed form solution to the minimizing solution with respect to the variational parameters [5, 55]. Minimizing this KL divergence is equivalent to maximizing the lower bound on the log likelihood (i.e., the evidence lower bound). Putting together Equations (12) and (7), we can optimize each of the variational parameters in closed form and conditionally on the others. In the fully factorized form, we can perform these updates sequentially; this produces a coordinate ascent optimization.

The optimal variational parameters ς_i that correspond to W_i are the solution to:

$$\log q_{\varsigma_i}(W_i) = E_{W_{-i}}[\log p(W, Y|X, \theta)] + C, \tag{13}$$

where the expectation is taken over all W except W_i according to the variational distribution.

The updates for the variational parameters, $\lambda_{k,1}$ and $\lambda_{k,2}$ of the stick breaking weights π are:

$$\lambda_{k,1} = \frac{\alpha}{K} + \sum_{q=1}^{Q} \eta_{q,k} \tag{14}$$

$$\lambda_{k,2} = 1 + \sum_{q=1}^{Q} (1 - \eta_{q,k}). \tag{15}$$

The updates for the variational parameters of the binary SNP-factor membership variables Z are:

$$\eta_{qk} = \frac{1}{1 + e^{-\zeta}} \tag{16}$$

$$
\begin{aligned}
\zeta &= (\psi(\lambda_{k1}) - \psi(\lambda_{k2})) \\
&- \frac{1}{2\sigma_n^2}\sum_{n=1}^{N}(X_{nq}(tr(\varphi_k) + \phi_{k.}\phi_{k.}^T)X_{nq}^T) \\
&- \frac{1}{\sigma_n^2}\sum_{n=1}^{N}(X_{n,q}(tr(\varphi_k) + \phi_{k.}\phi_{k.}^T)\eta_{q'k}^T X_{nq'}^T \\
&+ \frac{1}{\sigma_n^2}\sum_{n=1}^{N} X_{n,q}\phi_{k.}(Y_{n.}^T - \phi_{k'.}^T\eta_{.k'}^T X_{n.}^T)
\end{aligned} \tag{17}
$$

The updates for the variational parameters of effect sizes A are:

$$
\varphi_k = (\frac{\kappa_{k\cdot 1}}{\kappa_{k\cdot 2}} + \frac{1}{\sigma_n^2}\sum_{n=1}^{N}\sum_{q=1}^{Q} X_{n.}\eta_{.k}X_{n.}^T)^{-1} \tag{18}
$$

$$
\begin{aligned}
\phi_{k.} &= \sum_{n=1}^{N} X_{n.}\eta_{.k}[(Y_{n.})^T - \sum_{k'<k}(\phi_{k'.}^m)^T\eta_{.k'}^T X_{n.}^T) \\
&* (\frac{\kappa_{k\cdot 1}}{\kappa_{k\cdot 2}} + \frac{1}{\sigma_n^2}\sum_{n=1}^{N}\sum_{q=1}^{Q} X_{n.}\eta_{.k}X_{n.}^T)^{-1}
\end{aligned} \tag{19}
$$

Finally, the updates for σ_A^2 are:

$$
\kappa_{k\cdot 1} = c + \frac{1}{2} \tag{20}
$$

$$
\kappa_{k\cdot 2} = d + \sum_{k=1}^{K}\frac{\varphi_{kp} + \phi_{kp}^2}{2}. \tag{21}
$$

Full derivations of the variational updates are available in the Supplementary material.

To assess convergence, we used the following approach, inspired by the Geweke test statistic [26]. Let $m_{i,1}$ and $s_{i,1}$ and be the mean and the standard deviation, respectively, of the first 10% of the trace of ς_i (i.e., estimates of ς_i over n_1 iterations). Let $m_{i,2}$ and $s_{i,2}$ be the mean and standard deviation of the last 50% of the trace of ς_i (i.e., estimates of ς_i over n_2 iterations). Both percentages exclude 100 burn-in iterations. We assessed convergence every 100 iterations by checking for unequal variance for all ς_i, where the t-statistic t_i equals:

$$
t_i = \frac{m_{i1} - m_{i2}}{\sqrt{\frac{s_{i1}}{n_1} + \frac{s_{i2}}{n_2}}} \tag{22}
$$

P-values p_i are computed for each parameter type t_i. If $p_i \leq p_{thresh}$, where we set $p_{thresh} = 0.05$, the segments of that trace are significantly different and the model has not yet converged. We find that, in general, this variational approach converged within 500 iterations.

The algorithm for fitting the BERRRI model is included here (Algorithm 1).

Algorithm 1 BERRRI algorithm

procedure FIT BERRRI
  while $i <$*maxiter* and $p_i \leq p_{thresh}$ for at least one parameter **do**
    for $k = 1...K$ **do**
      Update $\lambda_{k,1}$ with Equation 14
      Update $\lambda_{k,2}$ with Equation 15
    end for
    for $k = 1...K$ **do**
      for $q = 1...Q$ **do**
        Update $\eta_{q,k}$ with Equation 16
      end for
    end for
    for $k = 1...K$ **do**
      Update φ_k with Equation 18
      Update ϕ_k with Equation 19
    end for
    for $k = 1...K$ **do**
      for $p = 1...P$ **do**
        Update $\kappa_{k,p,1}$ with Equation 20
        Update $\kappa_{k,p,2}$ with Equation 21
      end for
    end for
  end while
end procedure

Testing for eQTLs

To assess whether a SNP is an eQTL for a specific gene, we computed variational maximum a posteriori (VMAP) estimates for each SNP-gene pair. According to our model, a SNP is an eQTL for a gene if the SNP is in any factor K *and* factor K has a non-zero weight for the observed gene expression levels. We can interpret an estimate of variational parameter $\eta_{q,k}$ as the posterior probability that SNP q is included in factor k. We therefore compute the matrix of V_{MAP} estimates for SNP-gene pairs as follows:

$$V_{MAP} = E_q[Z]E_q[A] = \eta\phi \tag{23}$$

where $\mathbb{E}_q[A]$ and $\mathbb{E}_q[Z]$ are the expected value of parameters A and Z under the variational distributions.

To determine a level at which a V_{MAP} for a SNP-gene association should be considered significant, we performed permutation testing to identify a V_{MAP} threshold with a false discovery rate (FDR) of $\leq 10\%$.

Methods

Simulated Data

Simulated data were generated for $N = \{25, 100, 500\}$ individuals, $P = \{25\}$ genes, $Q = \{25, 100, 1000\}$ SNPs, and $K = \{5, 25\}$ latent factors. These data were generated using actual genotype data from HapMap Phase 3, chromosome 20. For each simulated data set, a random subset of N individuals and Q SNPs were selected from the genotype data. An inclusion matrix Z—indicating which SNPs are included in each factor—was simulated by, for each $k \in \{1, \ldots, K\}$, selecting a SNP at random to include in the factor, and then including additional SNPs based on the correlation of the genotypes across all individuals with that of the index SNP. A weight matrix A—encoding the effect size that a group of SNPs has on each of a subset of traits—was then simulated from a normal distribution with mean zero and standard deviation 0.5.

The gene expression matrix Y was then computed from the simulated parameters, and independent Gaussian random noise was added to each element with mean zero and standard deviation 1.

$$\epsilon_{i,j} \sim \mathcal{N}(0, 1) \tag{24}$$

$$Y = XAZ + \epsilon \tag{25}$$

The code to generate simulated data is included as Supplementary material.

Methods for comparison

We compared the results of BERRRI on these simulations with two other RRR methods: SRRR [9] and CRAM [17]. SRRR is a penalized least-squares RRR method for simultaneous dimension reduction and variable selection. Dimension reduction is accomplished by explicitly representing the Q by P matrix of regression coefficients as two rank K matrices. SRRR accomplishes variable selection with a group lasso

penalty [60], determining optimal parameters using the subgradient method [20].

CRAM uses an additive model to estimate each dimension P of the response, constraining the complexity of the model with the nuclear norm and encouraging the set of functions to be low rank [17]. The authors derive backfitting algorithms to estimate optimal parameters. Compared to RRR methods that use an orthogonal projection onto a lower rank space, the CRAM method is well defined for high-dimensional data.

SRRR requires two hyperparameters to be specified: λ, the penalty parameter, and K, the rank of the RRR. CRAM requires the specification of a single penalty parameter, λ. For both SRRR and CRAM, we estimated hyperparameters via five-fold cross-validation. Each method was subsequently run on the simulated data sets ten times using the optimal hyperparameters.

The computational complexity for both CRAM and SRRR is $O(NKQ^2)$ with some smaller linear terms, where Q is the dimension of the explanatory matrix, N the number of observations, and K the rank of the RRR. The computational complexity of BERRRI is $O(NPK^2Q^2)$ with some smaller linear terms, making clear the cost of the additional representation in terms of computational time.

Evaluation

For each method and each simulation, we measured the overall wall clock run time, residual sum of squares (RSS), and the precision and recall for the true associations recovered from the simulated data. The 95% confidence interval for each metric was computed.

HapMap Phase 3 eQTL study

We applied BERRRI to a subset of data from the HapMap Phase 3 study. These data include genotype data from 608 individuals from 11 worldwide populations [27]. Gene expression microarray data were collected on the lymphoblastoid cell lines (LCLs) from these 608 individuals [44]. We processed these microarray data to remove low expressed genes and technical and biological effects that may confound association [15]. We projected each of the gene expression levels to the quantiles of a standard normal to control for non-Gaussianity and outliers; the full data processing pipeline is described elsewhere [8].

We considered all expressed genes located on chromosome 21, and selected at random 1% of all non-redundant SNPs on the chromosome. We applied BERRRI to the gene expression matrix $Y \in \Re^{608x175}$ and genotype matrix $X \in \Re^{608x171}$. We ran the variational inference algorithm until convergence, and

used the point estimate of the parameters as our result. Finally, the variational map estimates for associations of each SNP with each gene were computed (Equation 23).

For comparison purposes, we also computed univariate Bayes factors for each SNP Q and gene P pair using BIMBAM (BF_U) [42]. Bayes factor 10% FDR thresholds for both approaches were identified via permutation testing.

Global FDR Thresholds

We determined the BF_U and VMAP thresholds for FDR $\leq 10\%$ using permuted data. More specifically, We fit the BERRRI model and computed the VMAP using (Equation 23) above for the true data, and then fit the BERRRI model and computed the VMAP again with permuted values of the response (i.e., shuffling sample labels of the gene expression matrix). We used the permuted VMAP estimates as a null VMAP distribution, assuming no associations. The same permuted data is used to compute a null distribution of BF_U.

Using these null distributions, we are able to compute a false discovery rate (FDR) for various thresholds of BF_U and VMAP. The FDR for a given threshold for each method is computed as the number of BFs or VMAPs exceeding the threshold in the permuted data (i.e., the number of false positives), divided by the number of BFs or VMAPs exceeding the threshold in the true data (i.e., the number of true positives plus false positives), scaled by the relative number of tests in the permuted and true data. For each method, we identified the threshold having less than or equal to 10% FDR to determine a cutoff for significant associations.

Discussion

Simulation results

To illustrate the effectiveness of the BERRRI approach, we simulated multiple correlated quantitative traits and genotype data in order to recover the associations among multiple SNPs and multiple traits. We compared our model to SRRR [9] and CRAM [17] (Figure 2 and Table 1).

BERRRI performs well relative to existing RRR methods CRAM and SRRR in the prediction of held out response variables (Figure 3; Table 1). SRRR outperforms BERRRI and CRAM for the RSS metric

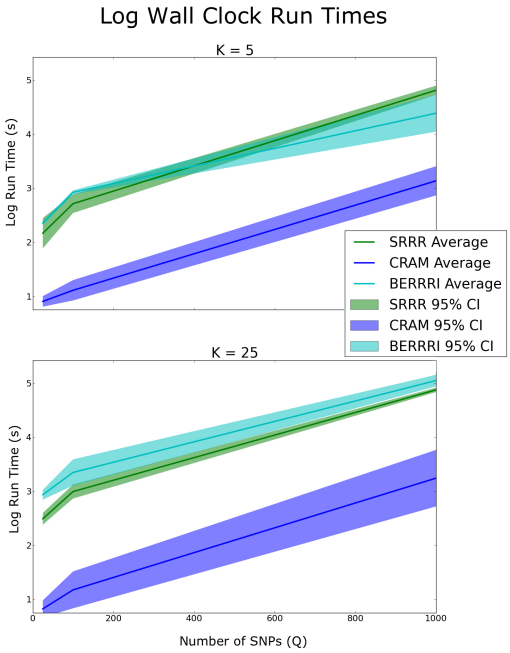

Figure 2: **Overall log wall clock run time in seconds for CRAM, SRRR, and BERRRI.** *x-axis: number of SNPs in the test; y-axis: log wall clock time. The dark lines show the average across ten runs; the ribbons show plus or minus three standard deviations from the mean run time.*

Table 1: Residual sum of squares (RSS) for the predicted trait values versus true response values in matrix Y for CRAM, SRRR, and BERRRI. Average and standard deviation are computed from ten runs. K represents the number of simulated factors.

Method	SNPs (Q)	True K	Avg. RSS	Std. Dev.
SRRR	25	5	1777.72	459.82
CRAM	25	5	**832.14**	1.9×10^{-3}
BERRRI	25	5	1921.60	103.54
SRRR	25	25	4,298.29	1,389.94
CRAM	25	25	4,160.82	8.9×10^{-3}
BERRRI	25	25	**3,343.82**	227.64
SRRR	100	5	4,101.31	1,246.95
CRAM	100	5	7,175.40	2.6×10^{-7}
BERRRI	100	5	**1,095.64**	100.02
SRRR	100	25	15,611.5	5,198.9
CRAM	100	25	41,331.5	1.2×10^{-6}
BERRRI	100	25	**2,707.07**	505.0
SRRR	1000	5	**97,798**	32,477
CRAM	1000	5	534,081	1.1×10^{-5}
BERRRI	1000	5	547,049	271,977
SRRR	1000	25	238,136	79,350
CRAM	1000	25	4,528,967	4.9×10^{-5}
BERRRI	1000	25	**157,692**	193,555

in one scenario when the simulated rank is low. This may due to SRRR requiring a prespecified rank, which we set to the true rank of the simulated data, giving SRRR an advantage.

The overall run time for BERRRI and SRRR are similar, with CRAM being the fastest. All scale approximately exponentially with the number of simulated predictors Q (Figure 2), which makes association testing across all SNPs infeasible without further approximations.

In order to quantify the proportion of true associations recovered using RRR models, we computed the precision-recall for each method based on the true associations in the simulated data (Figure 3). At high levels of recall (0.75 and greater), BERRRI shows dominant precision around 0.75; the other methods drop off in their precision much more sharply as the recall increases from 0.75 to 1.0. SRRR initially shows higher precision than the other methods, with a steady decline and falling below BERRRI near a 0.75 recall rate. For this analysis, we also ran SRRR with an incorrectly specified rank (SRRR_IR). As expected, this shows worse performance with a sharper decline in precision compared to SRRR. CRAM shows lower precision than the other methods for most recall rate levels. BERRRI recapitulates more of the underlying structure than SRRR and CRAM in the simulated data, while maintaining specificity and avoiding the detection of spurious associations.

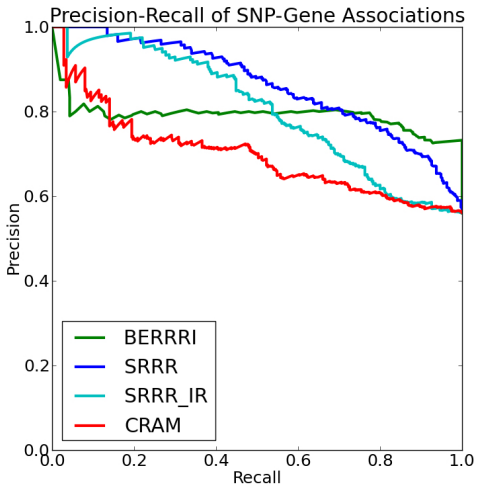

Figure 3: *Estimated precision and recall of CRAM, SRRR, SRRR with an incorrectly specified rank (SRRR_IR), and BERRRI. Precision and recall quantify the proportion of discovered associations that are truly significant and the proportions of truly significant associations that are discovered, respectively, quantified using simulated data with ground truth and the variational MAP estimates of effect size.*

Results on HapMap Phase 3 eQTL study data

To illustrate the effectiveness of the BERRRI approach in multi-SNP, multi-trait association mapping, we applied BERRRI to a subset of data from the HapMap Phase 3 Study [27]. We selected chromosome 21, choosing at random 1 % of all non-redundant SNPs, and all expressed genes on the chromosome. This resulted in a gene expression matrix $Y \in \Re^{608x175}$ and genotype matrix $X \in \Re^{608x171}$. We ran BERRRI on these data, and computed the BERRRI VMAP values to quantify associations of each SNP with each gene. We computed FDR using permutations.

To determine whether or not BERRRI identifies true associations, we compared the BERRRI $VMAP$ values for the SNP-gene associations with the univariate Bayes factors (BF_U) for each SNP-gene pair [42] (Figure 4). We observed more eQTL associations using BERRRI (n=157 at 10% FDR) than using the univariate approach (n=38 at 10% FDR). By analyzing the gene expression and SNP data jointly in BERRRI and by reducing the dimensionality of the genotype data via latent factors, we were able to obtain additional statistical power in association mapping. This additional power will lead to improvements in discovering eQTLs with smaller effect sizes relative to the univariate approach.

Importantly, BERRRI fails to discover some associations found to be significant with the univariate

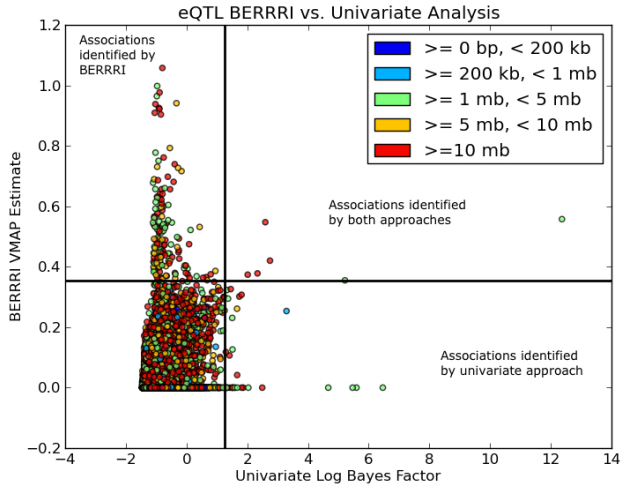

Figure 4: *BERRRI VMAP estimates for associations of each SNP with each gene from our model, versus Bayes factors from a classical eQTL univariate analysis [42] on Hapmap Phase 3 eQTL study data. The vertical line indicates the 10% FDR threshold for BERRRI, and the horizontal line indicates the 10% FDR for the classical univariate analyses. Points in the upper left quadrant are SNP-gene associations significant only in the BERRRI approach. Points in the upper right quadrant are associations significant in both approaches. Associations in the lower right quadrant are only significant in the classical univariate approach.*

analysis (Figure 4). It could be that these missing associations are either false negatives in the univariate analysis or missing (singleton) true associations. Statistically, this occurs because a SNP is excluded from a reduced rank association, which may happen when the association is explained by another highly correlated SNP.

We looked at the enrichment of proximal and distal associations identified by BERRRI and the univariate approach (Figure 4). To assess whether BERRRI identifies more distal associations than the classical approach, we performed a simple logistic regression, testing the hypothesis that distance between a SNP and a gene is significantly different for the associations identified by BERRRI only (upper left quadrant of Figure 4), and the associations identified by the univariate approach only (lower right quadrant). We found that there was a significant enrichment ($p = 0.0239$) for more distal associations in the eQTLs identified by BERRRI.

The most significant eQTL from the univariate approach is SNP rs2839146 for gene *YBEY*, which is located 75.7 kb upstream of the gene transcription start site (TSS); the BERRRI VMAP value for this association is 0.56 and the univariate log BF is 12.37 (Figure 5A). This eQTL is significant in both approaches, and was identified as statistically significant across twelve tissues in the Genotype-Tissue

Expression (GTEx) v6 data [3]. In the GTEx data, this SNP is also an eQTL for gene *MCM3AP*. It is located 24.5 kb away from the gene TSS, and was identified as statistically significant across thirteen tissues in GTEx. The BERRRI VMAP value for this association is 0.25 and the univariate log BF is 1.11 (Figure 5bB), both trending toward significance. BERRRI identified two other potential eQTLs for this gene: rs2839328 (329.4 kb away), and rs8133340 (331 kb away) with VMAP estimates of 0.52 and 0.43, respectively. Neither of these SNPs were listed as eQTLs for *MCM3AP* in GTEx, nor were they in LD with the GTEx listed eQTLs ($r^2 > 0.5$).

The second most significant eQTL from the univariate approach is SNP rs3992 for gene *PIGP*, which is located 74.5 kb from the gene TSS; the BERRRI VMAP value for this association is 0.0 and the univariate log BF is 6.47. This eQTL was identified as a statistically significant eQTL across twenty tissues in the GTEx data.

Another eQTL identified by both approaches is rs9974970 for gene *HLCS*, which is located 217.7 kb from the gene TSS; the BERRRI VMAP value for this association is 0.35 and the univariate log BF is 5.22. This eQTL was identified as statistically significant across sixteen tissues in the GTEx data [3].

BERRRI also identified eQTLs not discovered by the univariate approach. For example, SNP rs2839328 is identified as an eQTL for genes *FTCD* (428.3 kb), and *PRMT2* (40.5 kb), which were both significant eQTLs in the GTEx data. The BERRRI VMAP estimates are 0.53, and 0.61, respectively. This SNP is also identified by BERRRI as an eQTL for many other genes, including *ADAMTS1*, *LINC00160*, *RBM11*, *TRAPPC10*, and *USP16*, all of which have a BERRRI VMAP estimate higher than 0.90, and all of which were outside the GTEx cis-eQTL testing window of 1 Mb away from a gene TSS. The univariate approach was underpowered to detect these associations.

In many of these SNP-gene results, the associations discovered by BERRRI results in greater locus heterogeneity, meaning that the associated SNPs are spread out across a wider range of cis-loci; this behavior has been observed in other association models that similarly separate estimation of SNP association and effect size, as we do here [14, 15]. For *SH3BGR*, BERRRI identified an association found by the univariate approach with more significance, and identified one association that the univariate approach was underpowered to find. Similarly, with *ADAMTS5*, BERRRI identified an eQTL that went undetected by the univariate approach. Both of these SNP-gene pairs were outside the 1 Mb cis-eQTL testing window of GTEx.

For *ETS2* and *PCNT*, a few nominally significant associations detected by the univariate approach

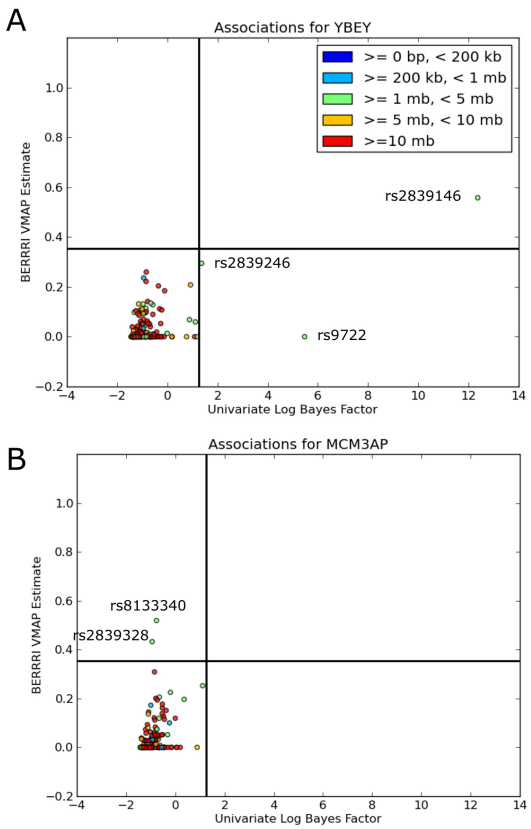

Figure 5: *The VMAP values from BERRRI and univariate approaches are shown on the vertical axes, with distance from the gene shown on the horizontal axis. The horizontal dashed lines correspond to the 10% FDR for each method. Panel A: Associations identified for YBEY. Panel B: Associations identified for MCM3AP.*

went undiscovered by BERRRI. We note that the undiscovered associations may either be a true association that is missed by BERRRI or a false positive univariate association that BERRRI avoids by sharing strength across genes and SNPs.

Imposing additional sparsity on the effect size matrix could further improve interpretation and association testing. For instance, a two-groups model such as a spike-and-slab Bayesian prior on A would result in an explicit posterior probability of association for eQTLs. This improvement may alleviate the challenge of detecting small effect eQTLs [14].

Conclusion

In this work, we developed a Bayesian nonparametric reduced rank regression method that effectively estimates a high-dimensional response matrix, inferring low-rank dimensionality from the data, and produces an interpretable model without any significant increase in computational time over state-of-the-art methods. While the results here illustrate an application to eQTL analyses, BERRRI may be applied to any multi-dimensional association analysis. We showed in simulations that our model accurately recovers latent structure in the response matrix and identifies associated predictors, improving over existing RRR methods lacking the interpretability and nonparametric behavior of the BERRRI approach. In the context of eQTL studies, BERRRI shows improved statistical power to identify genetic associations compared to classical univariate tests.

Acknowledgement

The authors would like to acknowledge Dr. Rina Foygel for providing CRAM software. This work was supported by the National Institutes of Health [R00 HG006265 and NIH R01 MH101822 to BEE]; and the Defense Advanced Research Projects Agency Predicting Health and Disease Award [N66001-09-C-2082 to GSG]. AMV was supported by the Gates Foundation [OPP1017554 to GSG].

References

[1] Theodore W Anderson and Herman Rubin. Estimation of the parameters of a single equation in a complete system of stochastic equations. *The Annals of Mathematical Statistics*, pp. 46–63, 1949.

[2] Theodore Wilbur Anderson. Estimating linear restrictions on regression coefficients for multivariate normal distributions. *The Annals of Mathematical Statistics*, pp. 327–351, 1951.

[3] Kristin G Ardlie, David S Deluca, Ayellet V Segrè, Timothy J Sullivan, Taylor R Young, Ellen T Gelfand, Casandra A Trowbridge, Julian B Maller, Taru Tukiainen, Monkol Lek, et al. The genotype-tissue expression (gtex) pilot analysis: Multitissue gene regulation in humans. *Science*, 348(6235):648–660, 2015.

[4] Hagai Attias. Inferring parameters and structure of latent variable models by variational Bayes. In *Proceedings of the Fifteenth conference on Uncertainty in artificial intelligence*, pp. 21–30. Morgan Kaufmann Publishers Inc., 1999.

[5] Matthew James Beal. *Variational algorithms for approximate Bayesian inference*. PhD thesis, University of London, 2003.

[6] David M Blei, Michael I Jordan, et al. Variational inference for dirichlet process mixtures. *Bayesian analysis*, 1(1):121–143, 2006.

[7] Rachel B Brem, Gaël Yvert, Rebecca Clinton, and Leonid Kruglyak. Genetic dissection of transcriptional regulation in budding yeast. *Science*, 296(5568):752–755, 2002.

[8] Christopher D Brown, Lara M Mangravite, and Barbara E Engelhardt. Integrative modeling of eQTLs and cis-regulatory elements suggests mechanisms underlying cell type specificity of eQTLs. *PLoS Genetics*, 9(8):e1003649, January 2013.

[9] Lisha Chen and Jianhua Z Huang. Sparse reduced-rank regression for simultaneous dimension reduction and variable selection. *Journal of the American Statistical Association*, 107(500):1533–1545, 2012.

[10] Vivian G Cheung and Richard S Spielman. Genetics of human gene expression: mapping dna variants that influence gene expression. *Nature Reviews Genetics*, 10(9):595–604, 2009.

[11] MuTHER Consortium et al. Identification of an imprinted master trans regulator at the *KLF14* locus related to multiple metabolic phenotypes. *Nature Genetics*, 43(6):561–564, 2011.

[12] William Cookson, Liming Liang, Gonçalo Abecasis, Miriam Moffatt, and Mark Lathrop. Mapping complex disease traits with global gene expression. *Nature Reviews Genetics*, 10(3):184–194, 2009.

[13] Finale Doshi, Kurt Tadayuki Miller, Jurgen Van Gael, and Yee Whye Teh. Variational inference for the indian buffet process. 2008.

[14] Bradley Efron et al. Microarrays, empirical bayes and the two-groups model. *Statistical Science*, 23(1):1–22, 2008.

[15] Barbara E Engelhardt and Ryan P Adams. Bayesian structured sparsity from Gaussian fields. *arXiv preprint arXiv:1407.2235*.

[16] Timothée Flutre, Xiaoquan Wen, Jonathan Pritchard, and Matthew Stephens. A statistical framework for joint eqtl analysis in multiple tissues. *PLoS genetics*, 9(5):e1003486, 2013.

[17] Rina Foygel, Michael Horrell, Mathias Drton, and John D Lafferty. Nonparametric reduced rank regression. In *Advances in Neural Information Processing Systems*, pp. 1628–1636, 2012.

[18] Nora Franceschini, Frank JA van Rooij, Bram P Prins, Mary F Feitosa, Mahir Karakas, John H Eckfeldt, Aaron R Folsom, Jeffrey Kopp, Ahmad Vaez, Jeanette S Andrews, et al. Discovery and fine mapping of serum protein loci through transethnic meta-analysis. *The American Journal of Human Genetics*, 91(4):744–753, 2012.

[19] Lude Franke and Ritsert C Jansen. eqtl analysis in humans. In *Cardiovascular Genomics*, pp. 311–328. Springer, 2009.

[20] Jerome Friedman, Trevor Hastie, Holger Höfling, Robert Tibshirani, et al. Pathwise coordinate optimization. *The Annals of Applied Statistics*, 1(2):302–332, 2007.

[21] Eric R Gamazon, Heather E Wheeler, Kaanan P Shah, Sahar V Mozaffari, Keston Aquino-Michaels, Robert J Carroll, Anne E Eyler, Joshua C Denny, Dan L Nicolae, Nancy J Cox, et al. A gene-based association method for mapping traits using reference transcriptome data. *Nature Genetics*, 47(9):1091–1098, 2015.

[22] C. Gao, C. D Brown, and B. E Engelhardt. A latent factor model with a mixture of sparse and dense factors to model gene expression data with confounding effects. *ArXiv e-prints*, October 2013.

[23] Chuan Gao, Christopher D Brown, and Barbara E Engelhardt. factors to model gene expression data with confounding effects. *arXiv preprint arXiv:1310.4792*, pp. 1–28.

[24] Chuan Gao, Shiwen Zhao, Ian C McDowell, and Barbara E Engelhardt. Differential gene co-expression networks via Bayesian biclustering models. *arXiv preprint arXiv:1411.1997*.

[25] John Geweke. Bayesian reduced rank regression in econometrics. *Journal of Econometrics*, 75(1):121–146, 1996.

[26] John Geweke et al. *Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments*, volume 196. Federal Reserve Bank of Minneapolis, Research Department Minneapolis, MN, USA, 1991.

[27] Richard A Gibbs, John W Belmont, Paul Hardenbol, Thomas D Willis, Fuli Yu, Huanming Yang, Lan-Yang Ch'ang, Wei Huang, Bin Liu, Yan Shen, et al. The international hapmap project. *Nature*, 426(6968):789–796, 2003.

[28] Thomas Griffiths and Zoubin Ghahramani. Infinite latent feature models and the Indian buffet process. 2005.

[29] Elin Grundberg, Kerrin S Small, Åsa K Hedman, Alexandra C Nica, Alfonso Buil, Sarah Keildson, Jordana T Bell, Tsun-Po Yang, Eshwar Meduri, Amy Barrett, et al. Mapping cis-and trans-regulatory effects across multiple tissues in twins. *Nature genetics*, 44(10):1084–1089, 2012.

[30] Elin Grundberg, Kerrin S Small, Åsa K Hedman, Alexandra C Nica, Alfonso Buil, Sarah Keildson, Jordana T Bell, Tsun-Po Yang, Eshwar Meduri, Amy Barrett, et al. Mapping cis-and trans-regulatory effects across multiple tissues in twins. *Nature genetics*, 44(10):1084–1089, 2012.

[31] Yongtao Guan and Matthew Stephens. Bayesian variable selection regression for genome-wide association studies and other large-scale problems. *The Annals of Applied Statistics*, pp. 1780–1815, 2011.

[32] Xin He, Chris K Fuller, Yi Song, Qingying Meng, Bin Zhang, Xia Yang, and Hao Li. Sherlock: detecting gene-disease associations by matching patterns of expression qtl and gwas. *The American Journal of Human Genetics*, 92(5):667–680, 2013.

[33] Qing Li, Ruibin Xi, Nan Lin, et al. Bayesian regularized quantile regression. *Bayesian Analysis*, 5(3):533–556, 2010.

[34] Zongming Ma and Tingni Sun. Adaptive sparse reduced-rank regression. *arXiv preprint arXiv:1403.1922*, 2014.

[35] Pekka Marttinen, Matti Pirinen, Antti-Pekka Sarin, Jussi Gillberg, Johannes Kettunen, Ida Surakka, Antti J Kangas, Pasi Soininen, Paul OReilly, Marika Kaakinen, et al. Assessing multivariate gene-metabolome associations with rare variants using bayesian reduced rank regression. *Bioinformatics*, p. btu140, 2014.

[36] Radford M. Neal. *Bayesian learning for neural networks*. PhD thesis, University of Toronto, 1995.

[37] Alexandra C Nica and Emmanouil T Dermitzakis. Expression quantitative trait loci: present and future. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 368(1620):20120362, 2013.

[38] Dan L Nicolae, Eric Gamazon, Wei Zhang, Shiwei Duan, M Eileen Dolan, and Nancy J Cox. Trait-associated snps are more likely to be eqtls: annotation to enhance discovery from gwas. *PLoS genetics*, 6(4):e1000888, 2010.

[39] VIVEK M Philip, Anna L Tyler, and Gregory W Carter. Dissection of complex gene expression using the combined analysis of pleiotropy and epistasis. In *Pacific Symposium on Biocomputing. Pacific Symposium on Biocomputing*, volume 19, p. 212. World Scientific, 2014.

[40] Joseph K Pickrell, John C Marioni, Athma A Pai, Jacob F Degner, Barbara E Engelhardt, Everlyne Nkadori, Jean-Baptiste Veyrieras, Matthew Stephens, Yoav Gilad, and Jonathan K Pritchard. Understanding mechanisms underlying human gene expression variation with rna sequencing. *Nature*, 464(7289):768–772, 2010.

[41] Sylvia Richardson, Leonardo Bottolo, and Jeffrey S Rosenthal. Bayesian models for sparse regression analysis of high dimensional data. *Bayesian Statistics*, 9:539–569, 2010.

[42] Bertrand Servin and Matthew Stephens. Imputation-based analysis of association studies: candidate regions and quantitative traits. *PLoS genetics*, 3(7):e114, 2007.

[43] Matthew Stephens. A unified framework for association analysis with multiple related phenotypes. *PLoS One*, 8(7):e65245, 2013.

[44] Barbara E Stranger, Stephen B Montgomery, Antigone S Dimas, Leopold Parts, Oliver Stegle, Catherine E Ingle, Magda Sekowska, George Davey Smith, David Evans, Maria Gutierrez-Arcelus, et al. Patterns of cis regulatory variation in diverse human populations. *PLoS Genetics*, 8(4):e1002639–e1002639, 2012.

[45] Barbara E Stranger, Eli A Stahl, and Towfique Raj. Progress and promise of genome-wide association studies for human complex trait genetics. *Genetics*, 187(2):367–383, 2011.

[46] Jae Hoon Sul, Buhm Han, Chun Ye, Ted Choi, and Eleazar Eskin. Effectively identifying eQTLs from multiple tissues by combining mixed model and meta-analytic approaches. *PLoS Genetics*, 9(6):e1003491, 2013.

[47] Yee W Teh, Dilan Görür, and Zoubin Ghahramani. Stick-breaking construction for the indian buffet process. In *International Conference on Artificial Intelligence and Statistics*, pp. 556–563, 2007.

[48] Yee W Teh, David Newman, and Max Welling. A collapsed variational bayesian inference algorithm for latent dirichlet allocation. In *Advances in neural information processing systems*, pp. 1353–1360, 2006.

[49] Romain Thibaux and Michael I Jordan. Hierarchical beta processes and the Indian buffet process. In *International conference on artificial intelligence and statistics*, pp. 564–571, 2007.

[50] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 267–288, 1996.

[51] Michael E. Tipping. Sparse Bayesian learning and the relevance vector machine. *The Journal of Machine Learning Research*, 1:211–244, 2001.

[52] William Valdar, Christopher C Holmes, Richard Mott, and Jonathan Flint. Mapping in structured populations by resample model averaging. *Genetics*, 182(4):1263–1277, 2009.

[53] Maria Vounou, Eva Janousova, Robin Wolz, Jason L Stein, Paul M Thompson, Daniel Rueckert, and Giovanni Montana. Sparse reduced-rank regression detects genetic associations with voxel-wise longitudinal phenotypes in alzheimer's disease. *Neuroimage*, 60(1):700–716, 2012.

[54] Maria Vounou, Thomas E Nichols, and Giovanni Montana. Discovering genetic associations with high-dimensional neuroimaging phenotypes: a sparse reduced-rank regression approach. *Neuroimage*, 53(3):1147–1159, 2010.

[55] Martin J Wainwright and Michael I Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1-2):1–305, 2008.

[56] M. West. Bayesian factor regression models in the "large p, small n" paradigm. In *Bayesian Statistics 7*, pp. 723–732. Oxford University Press, 2003.

[57] Melanie A Wilson, Edwin S Iversen, Merlise A Clyde, Scott C Schmidler, and Joellen M Schildkraut. Bayesian model search and multilevel inference for snp association studies. *The annals of applied statistics*, 4(3):1342, 2010.

[58] Patricia J Wittkopp, Belinda K Haerum, and Andrew G Clark. Regulatory changes underlying expression differences within and between drosophila species. *Nature genetics*, 40(3):346–350, 2008.

[59] Gregory A Wray. The evolutionary significance of cis-regulatory mutations. *Nature Reviews Genetics*, 8(3):206–216, 2007.

[60] Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(1):49–67, 2006.

[61] Wei Zhang, Jun Zhu, Eric E Schadt, and Jun S Liu. A bayesian partition method for detecting pleiotropic and epistatic eqtl modules. *PLoS computational biology*, 6(1):e1000642, 2010.

[62] Shiwen Zhao, Chuan Gao, Sayan Mukherjee, and Barbara E Engelhardt. Bayesian group latent factor analysis with structured sparse priors. *arXiv preprint arXiv:1411.2698v1*.

[63] Jun Zhu, Bin Zhang, Erin N Smith, Becky Drees, Rachel B Brem, Leonid Kruglyak, Roger E Bumgarner, and Eric E Schadt. Integrating large-scale functional genomic data to dissect the complexity of yeast regulatory networks. *Nature genetics*, 40(7):854–861, 2008.